Harvard John A. Paulson School of Engineering and Applied Sciences

The Concurrent Composition of Differential Privacy

Tianhao Wang, Princeton University

Master Thesis Project at Harvard University
Advisor: Prof. Salil Vadhan

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- Definitions and Basic Properties
- Concurrent Composition for Pure Interactive Differential Privacy
- Concurrent Composition for Approximate Interactive Differential Privacy
- Characterization of Concurrent Composition
- Empirical Findings & Future Work

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- **Background**
- Definitions and Basic Properties
- Concurrent Composition for Pure Interactive Differential Privacy
- Concurrent Composition for Approximate Interactive Differential Privacy
- Characterization of Concurrent Composition
- Empirical Findings & Future Work

Background: DP under Composition

- Goal: analyze the privacy loss under the composition of multiple different mechanisms on the same dataset
  - Rarely want to release only a single statistic about a dataset.
  - Useful tool in algorithm design.
  - If the building blocks are proven to be private, it would be easy to reason about privacy of a complex algorithm built on these building blocks.

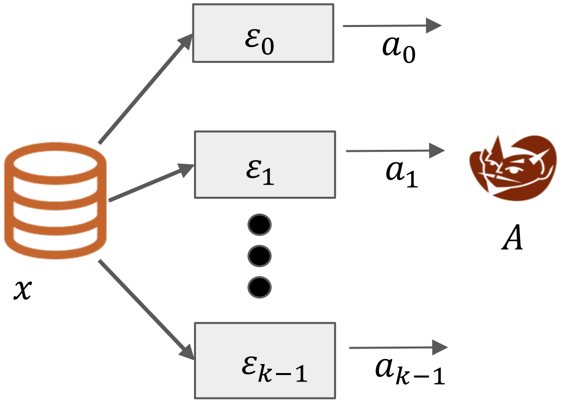

Background: DP under Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

- Basic composition: If M_i is (ε, δ) −DP for $i = 0, \ldots, k-1$, then $M(x) = (M_0(x), \ldots, M_{k-1}(x))$ is $(k\varepsilon, k\delta)$ −DP.
    - The randomness of M_i are independent to each other.
- Advanced composition: If M_i is (ε, δ) −DP for $i = 0, \ldots, k-1$, then $M(x) = (M_0(x), \ldots, M_{k-1}(x))$ is $(O(\sqrt{k \log(\frac{1}{\delta'})}\, \varepsilon, k\delta + \delta')$-DP.
    - The randomness of Mi are independent to each other.
- Optimal Composition [KOV15, MV16]
- Moment accountant [ACGMMTZ16]

Interactive Differential Privacy

- Many of the useful differential privacy primitives are actually interactive mechanisms, which allow one to ask an adaptive sequence of queries about the dataset.
  - e.g., **Sparse Vector Technique (SVT)**   Many applications!

$Input:\ q_1, \dots, q_i, \dots q_\infty$

$If\ q_i + \text{Noise} > \mathrm{T_i} + \text{Noise}, \text{output} \top;$

$else\ output\ \bot.$

$*\ c\ \#\ number\ of\ \top\ (\,Privacy\ cost\ is\ proportional\ to\ \sqrt{c}\,)$

(Adapted from Yuqing Zhu)

Interactive DP under Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

- There could be more than one composition operations for interactive mechanisms.
- **Sequential Composition**: all of the queries to the current mechanism must be completed before the interaction with another mechanism can be spawned.

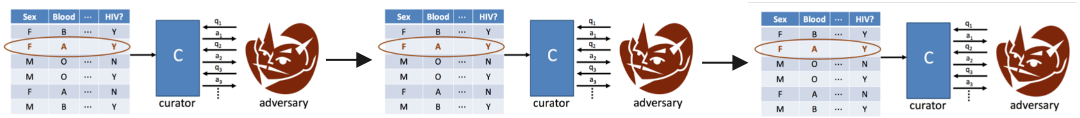

- **Concurrent Composition**: multiple interactions can be spawned and be executed simultaneously, queries to the mechanisms can be arbitrarily interleaved.

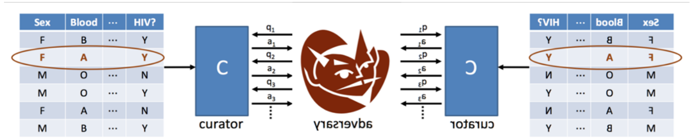

Main Results

Harvard John A. Paulson
School of Engineering
and Applied Sciences

- Group Privacy-like bound $(k\varepsilon, ke^{k\varepsilon}\delta)$ for the concurrent composition of approximate interactive DP mechanisms.

- Characterize arbitrary **pure** interactive DP mechanism as the interactive post-processing of randomized response (a non-interactive mechanism).
- **=> Optimal bound** for the concurrent composition of pure interactive DP.
- Based on computer simulation, we conjecture that optimal composition bound may extend to approximate DP.

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- **Definitions and Basic Properties**
- Concurrent Composition for Pure Interactive Differential Privacy
- Concurrent Composition for Approximate Interactive Differential Privacy
- Characterization of Concurrent Composition
- Empirical Findings & Future Work

Formalization: interactive protocol

- Interactive Protocol between two parties A and B
  - Each party as a function
  - (**private input**, **received messages**, **random coins**) => Next message to be sent out

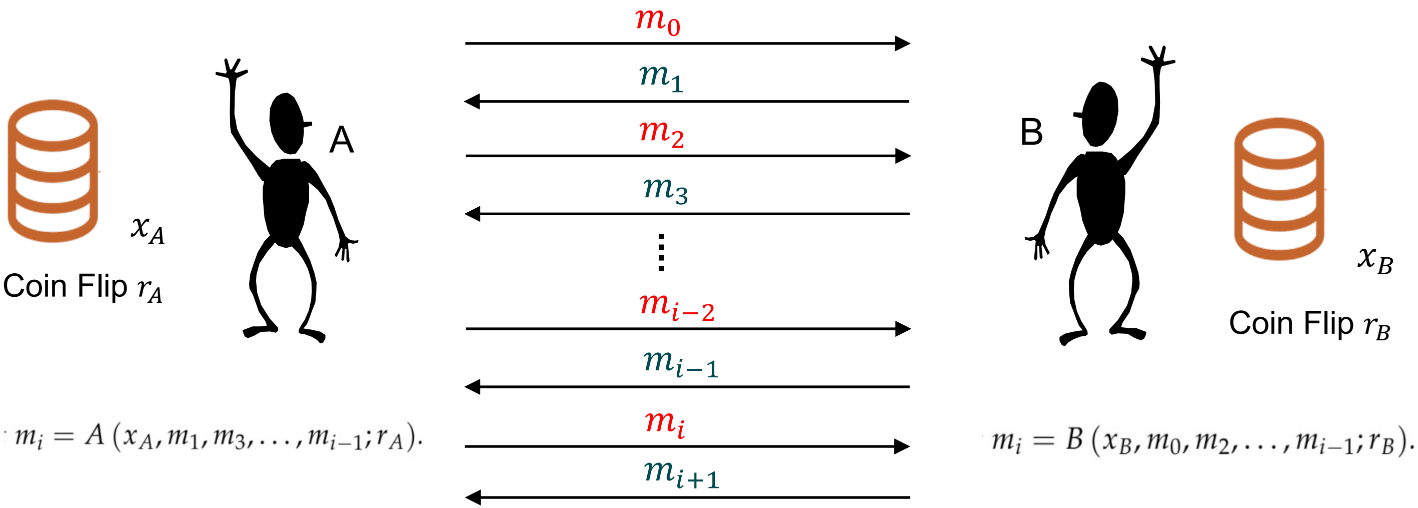

$$m_i = A\left(x_A, m_1, m_3, \ldots, m_{i-1}; r_A\right).$$

$$m_i = B\left(x_B, m_0, m_2, \ldots, m_{i-1}; r_B\right).$$

Formalization: view of a party

Harvard John A. Paulson School of Engineering and Applied Sciences

$$\texttt{View}_A\langle A(x_A), B(x_B)\rangle = (r_A, x_A, \underline{m_1, m_3, \ldots})$$

Randomness Private Input received messages

- In our case, party A is the adversary and party B is an interactive mechanism whose input is dataset x.
- Since we will only be interested in the adversary's view and the adversary does not have an input, we will drop the subscript and write A's view as $\texttt{View}\langle A, B(x)\rangle$

Formalization: interactive differential privacy

Harvard John A. Paulson School of Engineering and Applied Sciences

- The interactive differentially privacy as a type of interactive protocol between an adversary (without any computational limitations) and an interactive mechanism of special properties.

Definition 4 (Interactive Differential Privacy). *A randomized algorithm $\mathcal{M}$ is (ε, δ)-differentially private interactive mechanism if for every pair of adjacent datasets x, x', for every adversary algorithm $\mathcal{A}$, for every possible output set $T \subseteq \mathrm{Range}\left(\mathtt{View}\langle \mathcal{A}, \mathcal{M}(\cdot)\rangle\right)$ we have*

$$\Pr\left[\mathtt{View}\langle \mathcal{A}, \mathcal{M}(x)\rangle \in T\right] \leq e^{\varepsilon} \Pr\left[\mathtt{View}\langle \mathcal{A}, \mathcal{M}(x')\rangle \in T\right] + \delta$$

Concurrent Composition

- We use $ConComp(M_0, \ldots, M_{k-1})$ to denote the concurrently composed mechanism of $M_0, \ldots, M_{k-1}$.

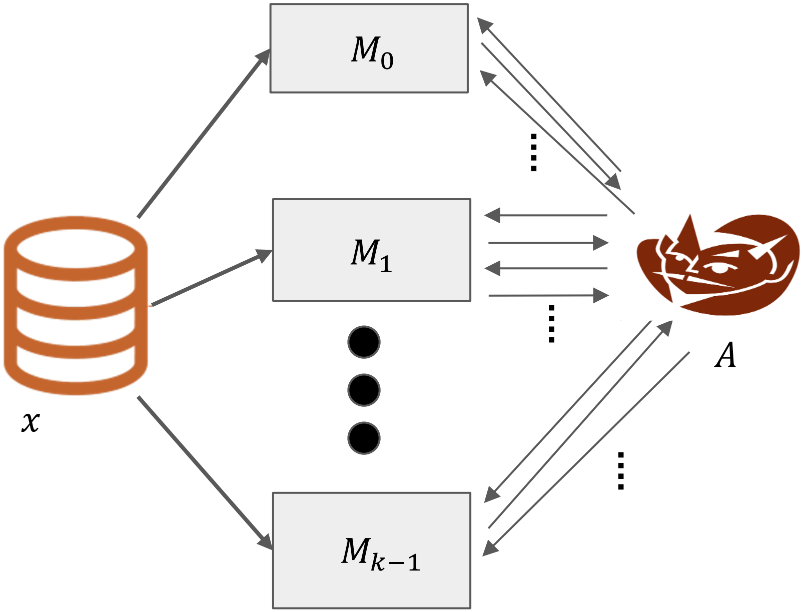

Ordered Concurrent Composition

- For the convenience of the proof, we introduce a variant of concurrent composition of interactive protocols, which only accept queries **in the exact order** of $M_0, \ldots, M_{k-1}$.

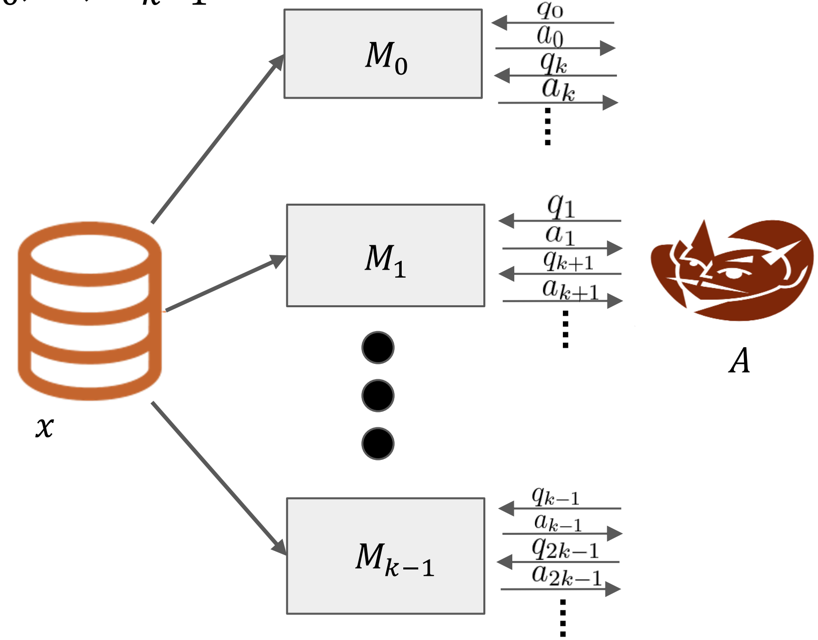

Ordered Concurrent Composition

- We also introduce the **null query extension** of an interactive mechanism, which has the exact same output distribution of the original mechanism but also accept “dummy” query strings.

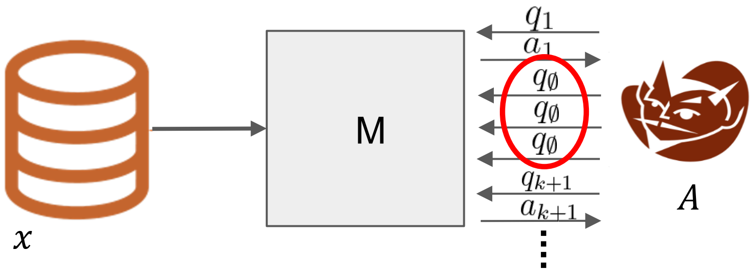

Ordered Concurrent Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

- Lemma: to prove $ConComp(M_0, \dots, M_{k-1})$ is (ε, δ)-DP, it suffices to prove the ordered concurrent composition of null query extensions of these mechanisms is also (ε, δ)-DP.
    - Intuition: if the first query is sent to M_i, the second query is not sent to M_{i+1} but M_j, we can simply fill in dummy queries between M_i and M_j.
- => We always assume the queries $q_0, \dots, q_{k-1}, q_k, \dots, q_{\{2k-1\}}, \dots$ from adversary are sent to $M_0, \dots, M_{k-1}$ **in order** in the proof, i.e., q_ℓ is sent to $M_{\ell \bmod k}$.
    - If an adversary A is concurrently interacting with two mechanisms M_0, M_1, we assume the queries **alternates** between them.

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- Definitions and Basic Properties
- **Concurrent Composition for Pure Interactive Differential Privacy**
- Concurrent Composition for Approximate Interactive Differential Privacy
- Characterization of Concurrent Composition
- Empirical Findings & Future Work

Concurrent Composition of Pure Interactive Differential Privacy

Harvard John A. Paulson School of Engineering and Applied Sciences

- Proof Sketch: view A and M as a combined adversary A^* interacting with $\tilde{M}$

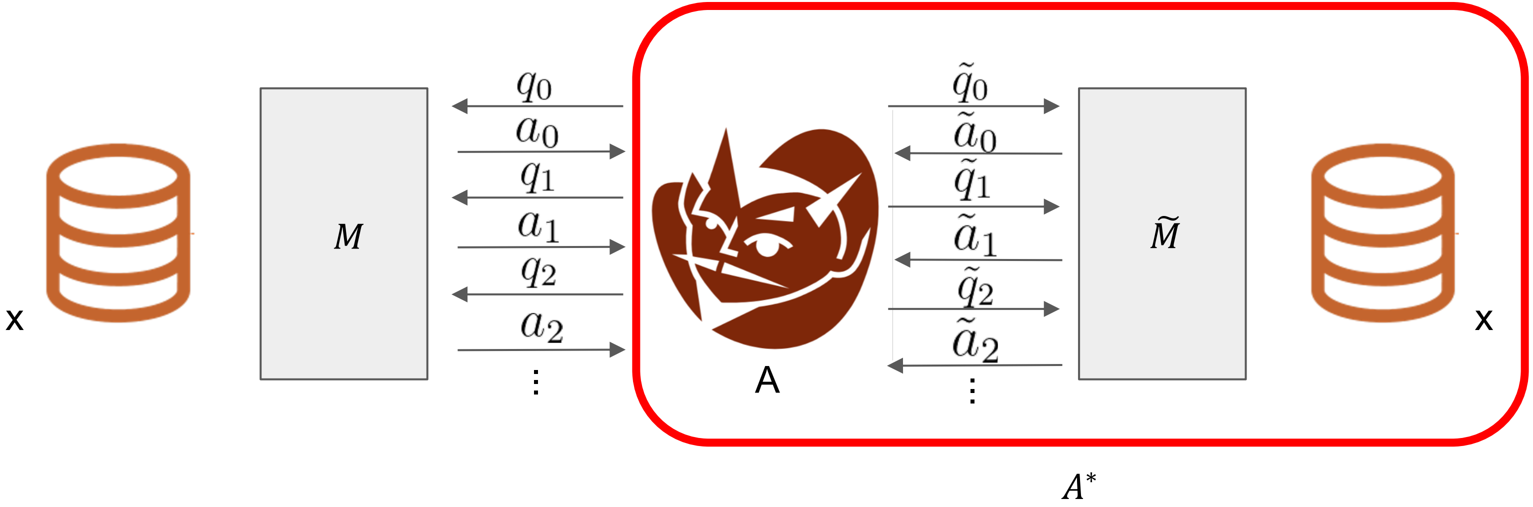

Concurrent Composition of Pure Interactive Differential Privacy

- Proof Sketch: view A and M as a combined adversary A^* interacting with $\widetilde{M}$
  - A^* is a well-defined strategy throughout the entire interactive session with M
    - Randomness of A^*: the randomness of A and $\widetilde{M}$
    - Next-query function is also naturally defined:

1. Random coin tosses for $\mathcal{A}^*_{\tilde{\mathcal{M}}}(x)$ consist of $r = (r_{\mathcal{A}}, r_{\tilde{\mathcal{M}}})$.
2. $\mathcal{A}^*_{\tilde{\mathcal{M}}}(x)(a_0, a_1, \ldots, a_{i-1}; r)$ is computed as follows:
   (a) $\tilde{q}_{i-1} = \mathcal{A}(a_0, \tilde{a}_0, \ldots, a_{i-1}; r_{\mathcal{A}})$, send to $\tilde{\mathcal{M}}$.
   (b) $\tilde{a}_{i-1} = \tilde{\mathcal{M}}(x, \tilde{q}_0, \tilde{q}_1, \ldots, \tilde{q}_{i-1}; r_{\tilde{\mathcal{M}}})$, send to $\mathcal{A}$.
   (c) $q_i = \mathcal{A}(a_0, \tilde{a}_0, \ldots, a_{i-1}, \tilde{a}_{i-1}; r_{\mathcal{A}})$.
   (d) Output q_i.

Concurrent Composition of Pure Interactive Differential Privacy

- Proof Sketch: view A and M as a combined adversary A^* interacting with $\widetilde{M}$
  - Given a transcript of A^*'s view, we can recover the view of A through post-processing, which is formulated as follows:

$$\texttt{Post}\left(r_{\mathcal{A}}, r_{\tilde{\mathcal{M}}}, a_0, a_1, \ldots, a_{T-1}; \mathcal{A}, \tilde{\mathcal{M}}(x)\right):$$

1. For $i = 1 \ldots T-1$, compute
   (a) $\tilde{q}_{i-1} = \mathcal{A}(a_0, \tilde{a}_0, \ldots, a_{i-1}; r_{\mathcal{A}})$
   (b) $\tilde{a}_{i-1} = \tilde{\mathcal{M}}(x, \tilde{q}_1, \tilde{q}_2, \ldots, \tilde{q}_{i-1}; r_{\tilde{\mathcal{M}}})$
2. Output $(r_{\mathcal{A}}, a_0, \tilde{a}_0, \ldots, a_{T-1}, \tilde{a}_{T-1})$.

Concurrent Composition of Pure Interactive Differential Privacy

Harvard John A. Paulson School of Engineering and Applied Sciences

- Proof Sketch: view A and M as a combined adversary A^* interacting with $\widetilde{M}$
  - Given a transcript of A^*'s view, we can recover the view of A through post-processing.
  - For any event T, the probability that A's view is in T is exactly equal to the probability that A^*'s view is in the inverse of the post-processing algorithm of event T.

$$\Pr\left[\mathtt{View}\langle\mathcal{A},\mathrm{ConComp}(\mathcal{M}(x),\tilde{\mathcal{M}}(x))\rangle \in T\right] = \Pr\left[\mathtt{View}\langle\mathcal{A}^*_{\tilde{\mathcal{M}}}(x),\mathcal{M}(x)\rangle \in \mathtt{Post}^{-1}(T)\right]$$

Concurrent Composition of Pure Interactive Differential Privacy

- Proof Sketch: view A and M as a combined adversary A^* interacting with $\widetilde{M}$
  - Given a transcript of A^*'s view, we can recover the view of A through post-processing.
  - For any event T, the probability that A's view is in T is exactly equal to the probability that A^*'s view is in the inverse of the post-processing algorithm of event T.
  - The random variable of A^*'s view enjoys differential privacy guarantee.

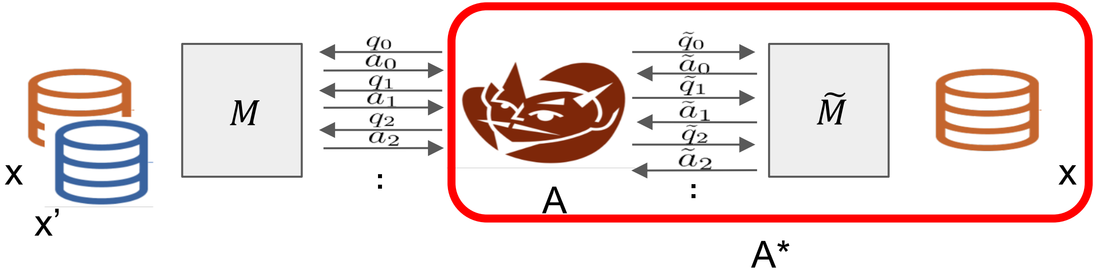

Concurrent Composition of Pure Interactive Differential Privacy

- We say two random variables X and X' are (ε, δ)-indistinguishable if for every event T we have

$$\Pr[X \in T] \leq e^{\varepsilon} \cdot \Pr[X' \in T] + \delta$$

$$\Pr[X' \in T] \leq e^{\varepsilon} \cdot \Pr[X \in T] + \delta$$

- Denote as $X \overset{(\epsilon,\delta)}{\approx} X'$
- Simple property: if $X \overset{(\epsilon,0)}{\approx} X'$, and $X' \overset{(\tilde{\epsilon},0)}{\approx} X''$, then $X \overset{(\varepsilon+\tilde{\varepsilon},0)}{\approx} X''$

Concurrent Composition of Pure Interactive Differential Privacy

Harvard John A. Paulson School of Engineering and Applied Sciences

- Suppose M is $(\varepsilon, 0)$-DP, $\tilde{M}$ is $(\tilde{\varepsilon}, 0)$-DP
- We know that $\texttt{View}(A^*_{\tilde{M}(x)}, M(x))$ and $\texttt{View}(A^*_{\tilde{M}(x)}, M(x'))$ are $(\varepsilon, 0)$-indistinguishable.
- => $\texttt{View}(A, \mathrm{ConComp}(M(x), \tilde{M}(x)))$ and $\texttt{View}(A, \mathrm{ConComp}(M(x'), \tilde{M}(x)))$ are $(\varepsilon, 0)$-indistinguishable.
- Symmetrically, we have $\texttt{View}(A, \mathrm{ConComp}(M(x'), \tilde{M}(x))) \overset{(\tilde{\varepsilon},0)}{\approx} \texttt{View}(A, \mathrm{ConComp}(M(x'), \tilde{M}(x')))$
- Finally, we can bound the privacy of A's view in the concurrent composition when the underlying dataset is x vs x'.

$$\texttt{View}(A, \mathrm{ConComp}(M(x), \tilde{M}(x))) \overset{(\varepsilon+\tilde{\varepsilon},0)}{\approx} \texttt{View}(A, \mathrm{ConComp}(M(x'), \tilde{M}(x')))$$

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- Definitions and Basic Properties
- Concurrent Composition for Pure Interactive Differential Privacy
- **Concurrent Composition for Approximate Interactive Differential Privacy**
- Characterization of Concurrent Composition
- Empirical Findings & Future Work

Concurrent Composition of Approximate Interactive Differential Privacy

- Suppose interactive mechanisms $M_0, \dots, M_{\{k-1\}}$ are each (ϵ_i, δ_i)-differentially private.
- View A and $M_0, \dots, M_{\{i-1\}}, M_{\{i+1\}}, \dots, M_{\{k-1\}}$ as a combined adversary A^*, we can show that:

$$\Pr\left[\mathtt{View}\langle \mathcal{A}, \mathrm{ConComp}(\mathcal{M}_0(x'), \dots, \mathcal{M}_{i-1}(x'), \mathcal{M}_i(x), \dots, \mathcal{M}_{k-1}(x))\rangle \in S\right]$$

$$\leq e^{\varepsilon_i} \Pr\left[\mathtt{View}\langle \mathcal{A}, \mathrm{ConComp}(\mathcal{M}_0(x'), \dots, \mathcal{M}_{i-1}(x'), \mathcal{M}_i(x'), \dots, \mathcal{M}_{k-1}(x))\rangle \in S\right] + \delta_i$$

Group Privacy-like Bound

$$\Pr\left[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x), \mathcal{M}_1(x), \ldots, \mathcal{M}_{k-1}(x))\rangle \in S\right]$$

$$\leq e^{\varepsilon_0} \Pr\left[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x'), \mathcal{M}_1(x), \ldots, \mathcal{M}_{k-1}(x))\rangle \in S\right] + \delta_0$$

$$\leq e^{\varepsilon_0}(e^{\varepsilon_1} \Pr\left[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x'), \mathcal{M}_1(x'), \ldots, \mathcal{M}_{k-1}(x))\rangle \in S\right] + \delta_1) + \delta_0$$

$$\leq \ldots$$

$$\leq e^{\sum_{i=0}^{k-1} \varepsilon_i} \Pr\left[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x'), \mathcal{M}_1(x'), \ldots, \mathcal{M}_{k-1}(x'))\rangle \in S\right]$$

$$+ \left(\delta_0 + e^{\varepsilon_0}\delta_1 + e^{\varepsilon_0+\varepsilon_1}\delta_2 + \ldots + e^{\sum_{i=0}^{k-2} \varepsilon_i}\delta_{k-1}\right) \leq k e^{\sum_{i=0}^{k-1} \varepsilon_i} \max_i(\delta_i)$$

$$= \left(1 + e^{\varepsilon} + e^{2\varepsilon} + \cdots + e^{(k-1)\cdot\varepsilon}\right) \cdot \delta$$

Same bound for Group Privacy

Group Privacy-like Bound

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- Definitions and Basic Properties
- Concurrent Composition for Pure Interactive Differential Privacy
- Concurrent Composition for Approximate Interactive Differential Privacy
- **Characterization of Concurrent Composition**
- Empirical Findings & Future Work

Characterization of Concurrent Composition

- Randomized Response

$$\mathrm{RR}_{(\varepsilon,\delta)} : \{0,1\} \to \{0,1, \texttt{`Iam0'}, \texttt{`Iam1'}\}$$ is (ε, δ)-DP

$$\Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(0) = \texttt{`Iam0'}\right] = \delta \qquad \Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(1) = \texttt{`Iam0'}\right] = 0$$

$$\Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(0) = 0\right] = (1-\delta) \cdot \frac{e^\varepsilon}{1+e^\varepsilon} \qquad \Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(1) = 0\right] = (1-\delta) \cdot \frac{1}{1+e^\varepsilon}$$

$$\Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(0) = 1\right] = (1-\delta) \cdot \frac{1}{1+e^\varepsilon} \qquad \Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(1) = 1\right] = (1-\delta) \cdot \frac{e^\varepsilon}{1+e^\varepsilon}$$

$$\Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(0) = \texttt{`Iam1'}\right] = 0 \qquad \Pr\left[\mathrm{RR}_{(\varepsilon,\delta)}(1) = \texttt{`Iam1'}\right] = \delta$$

[KOV15, MV17]

Characterization of Concurrent Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

- For **every** non-interactive (ε, δ) $-$DP algorithms and every neighboring dataset $x_0 \sim x_1$, there exists a post-processing T of randomized response $RR_{(\varepsilon,\delta)}$ such that $T(RR_{(\varepsilon,\delta)}(b))$ is identically distributed to $M(x_b)$ [KOV15, MV17].

- Post-processing preserves differential privacy

  =>To analyze the composition of arbitrary non-interactive DP algorithms, it suffices to analyze the composition of RR's.

- If we are able to prove a similar result for interactive differential privacy, then we will be able to extend all results of composition theorem for non-interactive mechanisms to interactive mechanisms!

Characterization of Concurrent Composition

- **Every** interactive $(\varepsilon, 0)$ –DP mechanisms can be simulated as the post-processing of randomized response $RR_{(\varepsilon,0)}$.

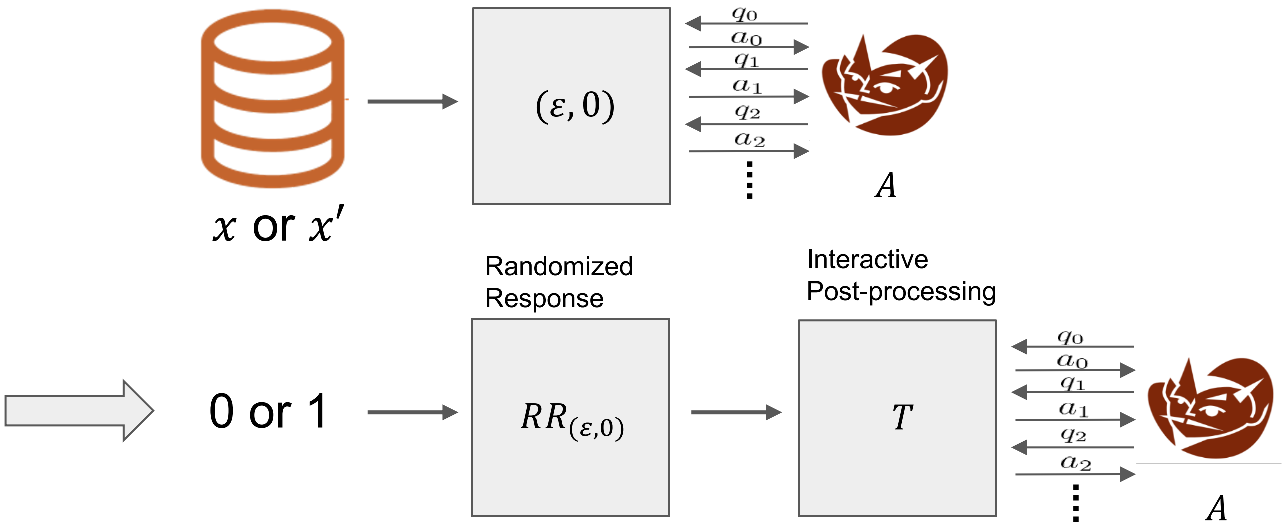

Characterization of Concurrent Composition

- **Every** interactive $(\varepsilon, 0)$ –DP mechanisms can be simulated as the post-processing of randomized response $RR_{(\varepsilon,0)}$.

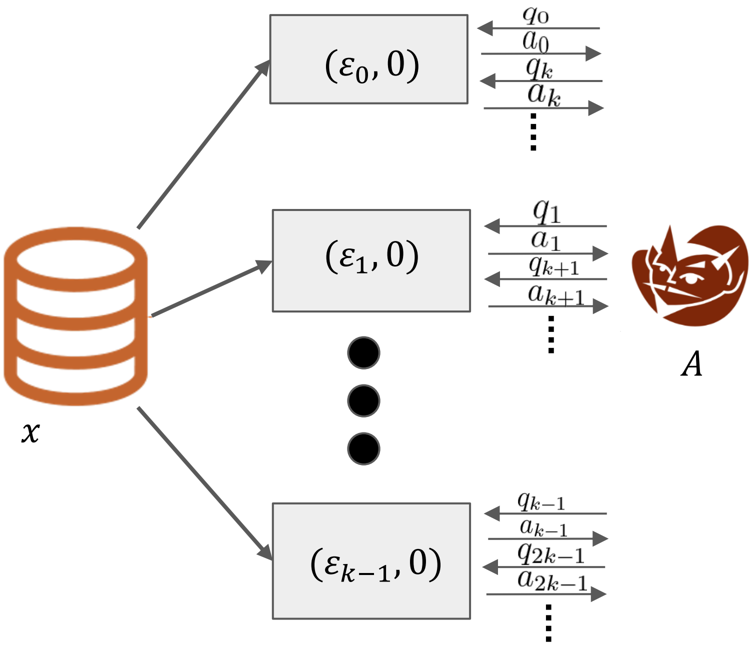

Characterization of Concurrent Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

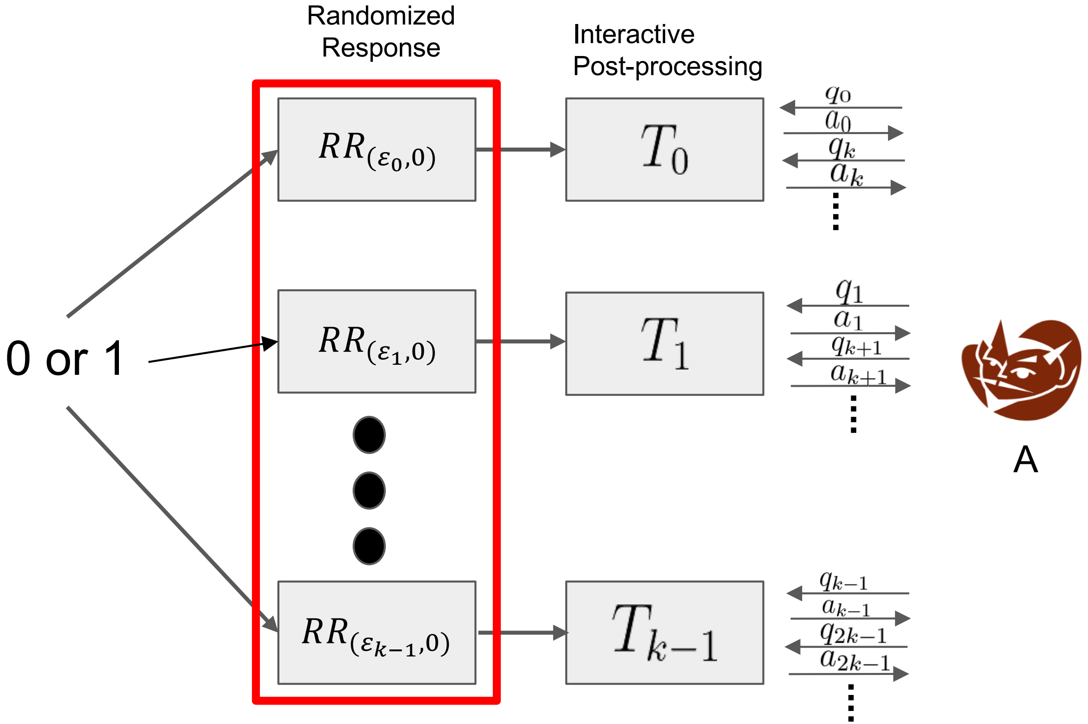

Characterization of Concurrent Composition

Harvard John A. Paulson School of Engineering and Applied Sciences

- If interactive mechanism $M_0, \ldots, M_{k-1}$ are each $(\varepsilon_i, 0)$-DP for $i = 0..k-1$, then given a target δ_g, the privacy parameter of the current composition $ConComp(M_0, \ldots, M_{k-1})$ is tightly upper bounded by the least value of ε_g such that

$$\frac{1}{\prod_{i=0}^{k-1}(1+e^{\varepsilon_i})} \sum_{S \subseteq \{0,\ldots,k-1\}} \max\left\{e^{\sum_{i \in S} \varepsilon_i} - e^{\varepsilon_g} \cdot e^{\sum_{i \notin S} \varepsilon_i}, 0\right\} \leq \delta_g$$

(Optimal Bound from MV17 for non-interactive DP)

Outline

Harvard John A. Paulson School of Engineering and Applied Sciences

- Background
- Definitions and Basic Properties
- Concurrent Composition for Pure Interactive Differential Privacy
- Concurrent Composition for Approximate Interactive Differential Privacy
- Characterization of Concurrent Composition
- **Empirical Findings & Future Work**

Empirical Findings & Future Work

- We find empirical evidence supports that the Optimal Composition Theorems from [KOV15] can be extended to the concurrent composition of approximate DP mechanisms.
  - We evaluate whether any 2-round (ϵ, δ) interactive mechanisms with 1-bit messages can be simulated by some interactive post-processing of randomized response.